NERSC – National Energy Research Scientific Computing Center

Presentation to the NRC Panel on "The Future of Supercomputing"

Horst D. Simon

Director, NERSC Center Division, LBNL

Washington D.C., December 4, 2003

http://www.nersc.gov/~simon

Flagship facility for unclassified computing in the DOE Office of Science

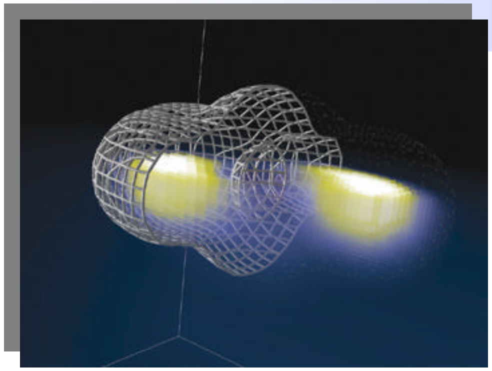

- computational projects in all areas of DOE science
- about 2000 users in ~200 projects

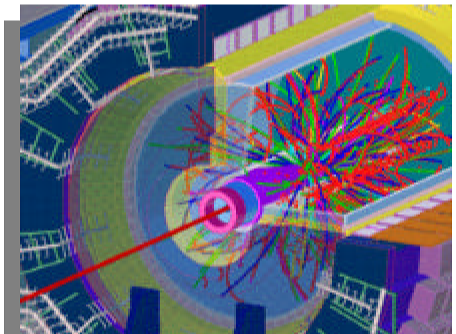

- focus on large-scale computing

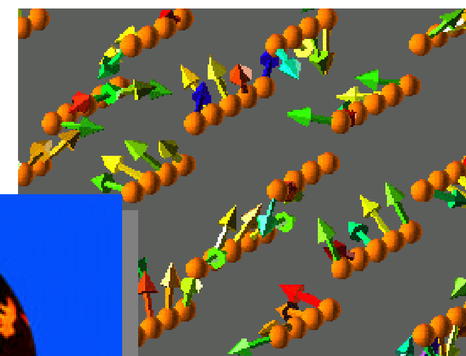

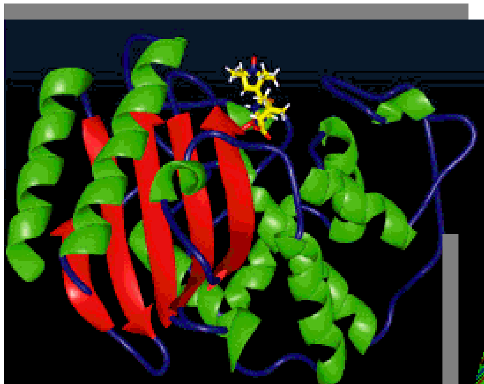

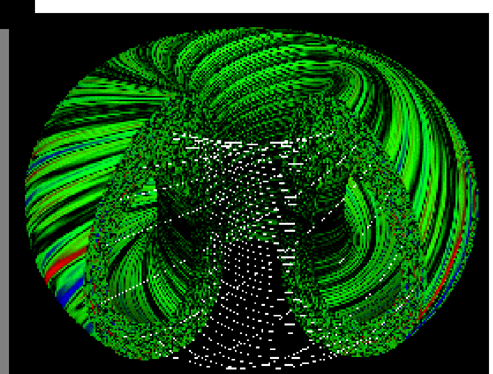

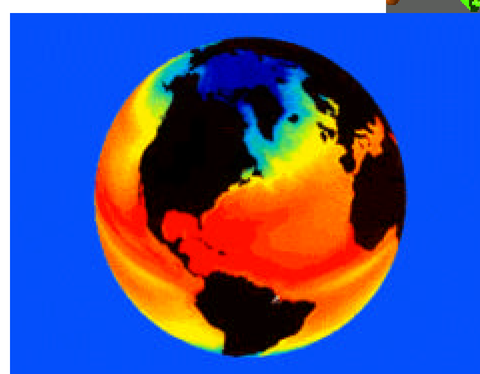

Capability Computing at NERSC

- Capability means
  - Large amount of time
  - Large amount of memory
  - Large number of processors
  - Large amount of data
- Focus on large scale computational science that cannot be done elsewhere

NERSC Center Overview

- Funded by DOE, annual budget $28M, about 65 staff
  - Traditional strategy to invest equally in newest compute platform, staff, and other resources
- Supports open, unclassified, basic research
- Close collaborations between university and NERSC in computer science and computational science

Office of Science
U.S. DEPARTMENT OF ENERGY

Components of the Next-Generation NERSC

New Technology Introduction

- About 25% of the staff dedicated to innovation and evolving services
  - Large data resources and management
  - Grid access
  - Innovative algorithms, software, and tools
  - New systems (computing, storage, and visualization)
- Example: Global Unified Parallel File System
  - Leading activity in addressing a common requirement for large scale shared storage

Usage by Discipline and Institution Type

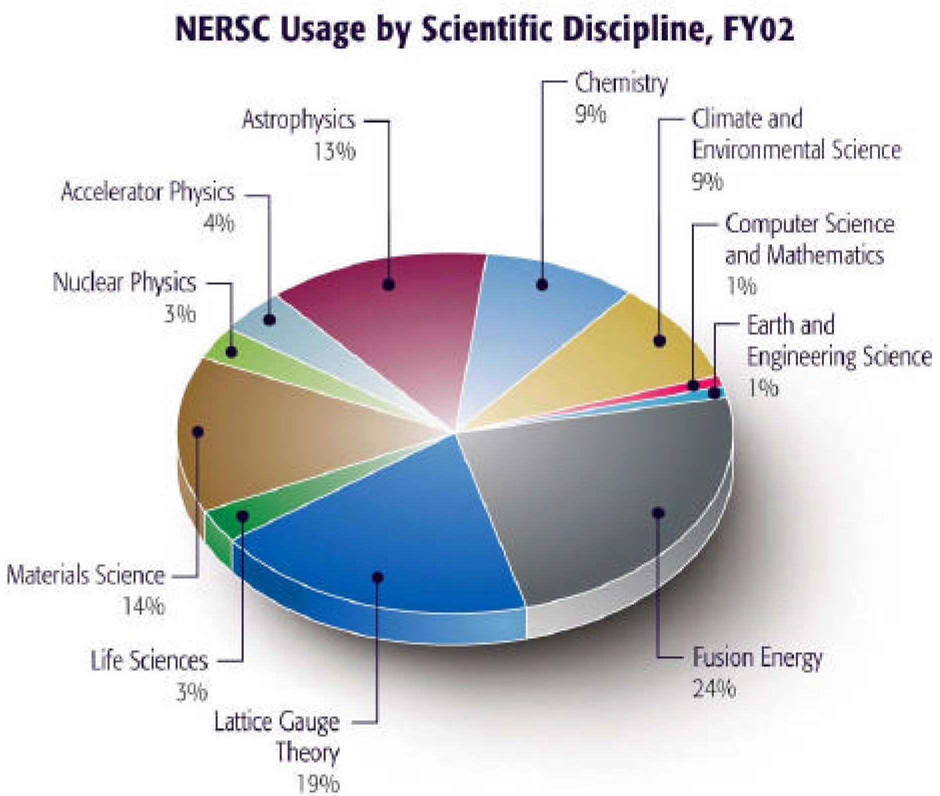

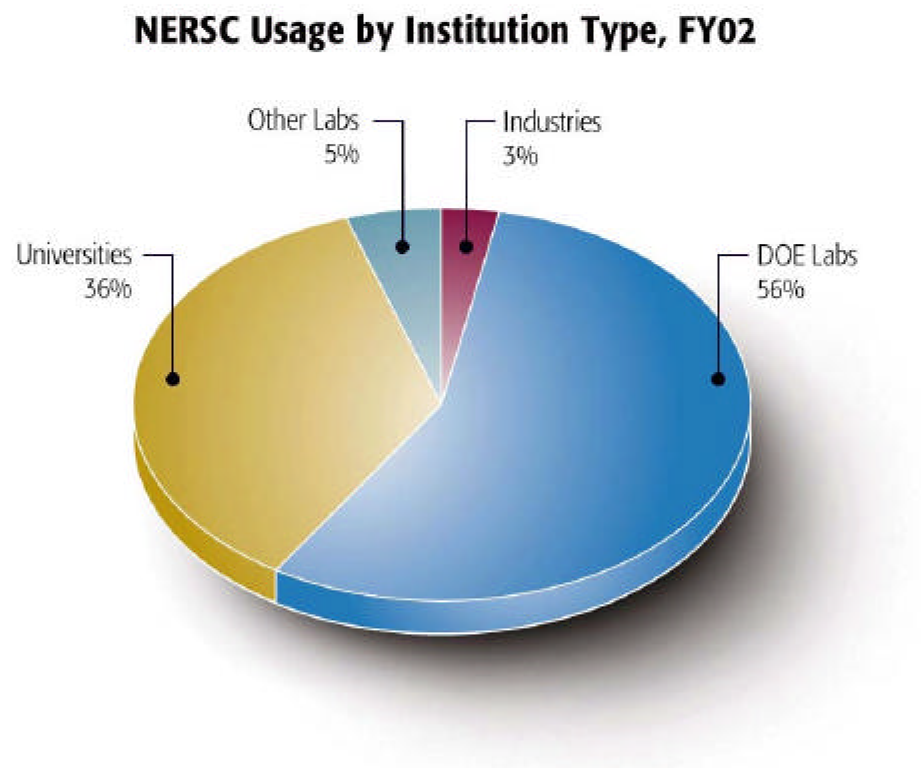

Usage by PI Institution

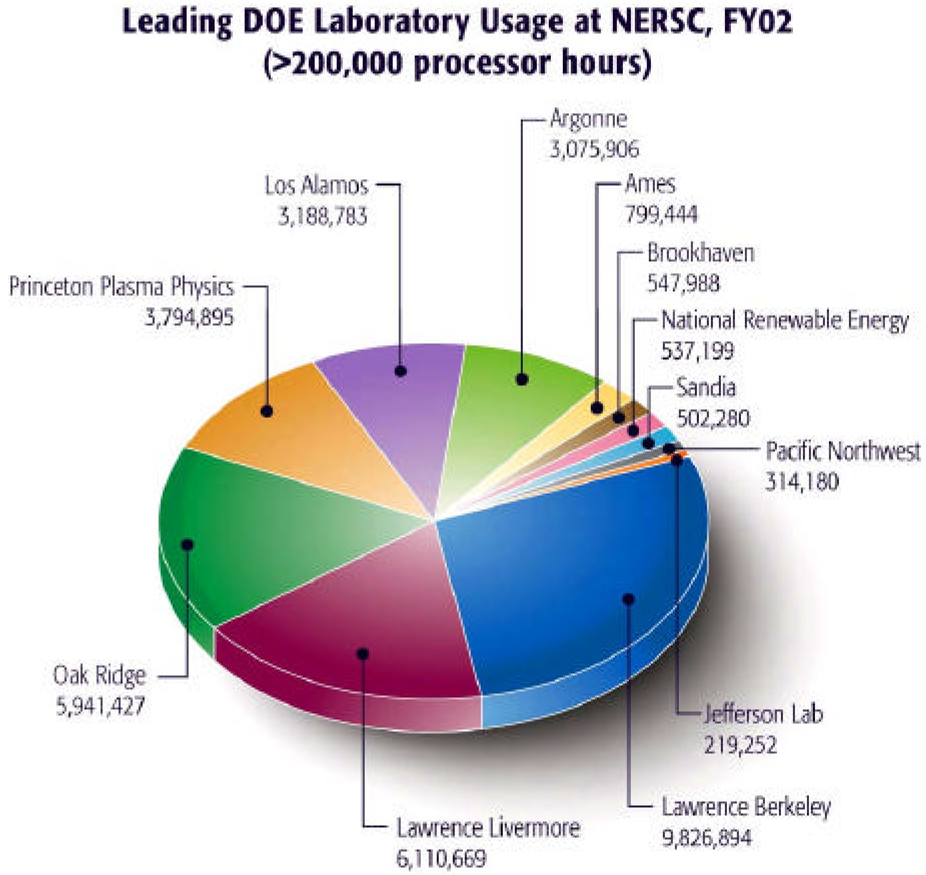

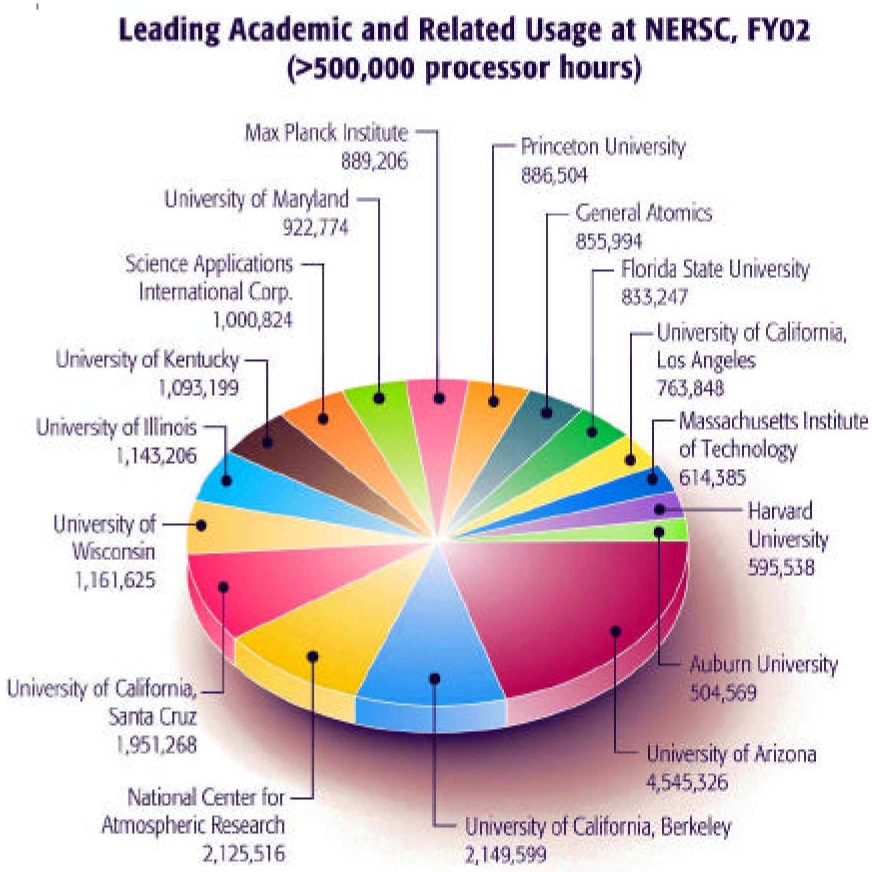

“Seaborg” Characteristics

- 416 16-way Power 3+ nodes with each CPU at 1.5 Gflop/s
- 6,656 CPUs – 6,080 for computation
- Total Peak Performance of 10 Tflop/s
- Total Aggregate Memory is 7.8 TB
- Total GPFS disk will be 44 TB

- One of the largest computational platforms in production use

NERSC System Architecture

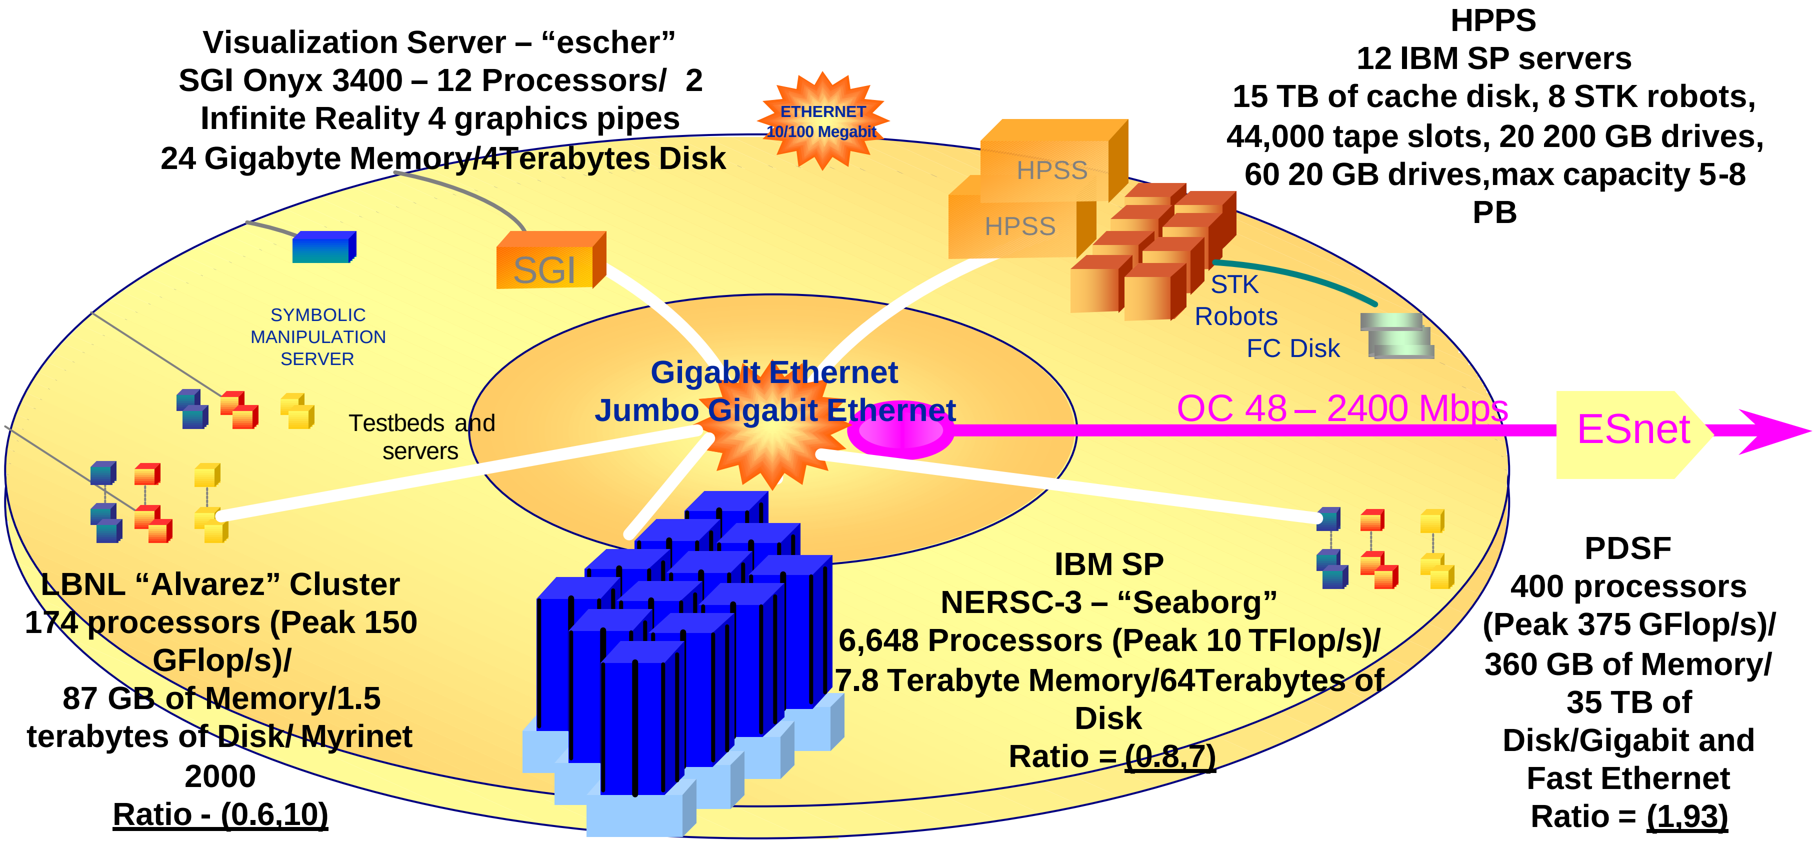

Ratio = (RAM Bytes per Flop, Disk Bytes per Flop)

Office of Science
U.S. DEPARTMENT OF ENERGY

Capability Computing in a Production Environment

- Seaborg has reached 95% utilization
- Large jobs (> 512 processors) account for about 50% of usage
- Number of projects declined from more than 400 in 2001 to less than 200 in 2003
- INCITE – special allocation for 1- 3 projects using 10% of Seaborg

Large Job Sizes Run Regularly

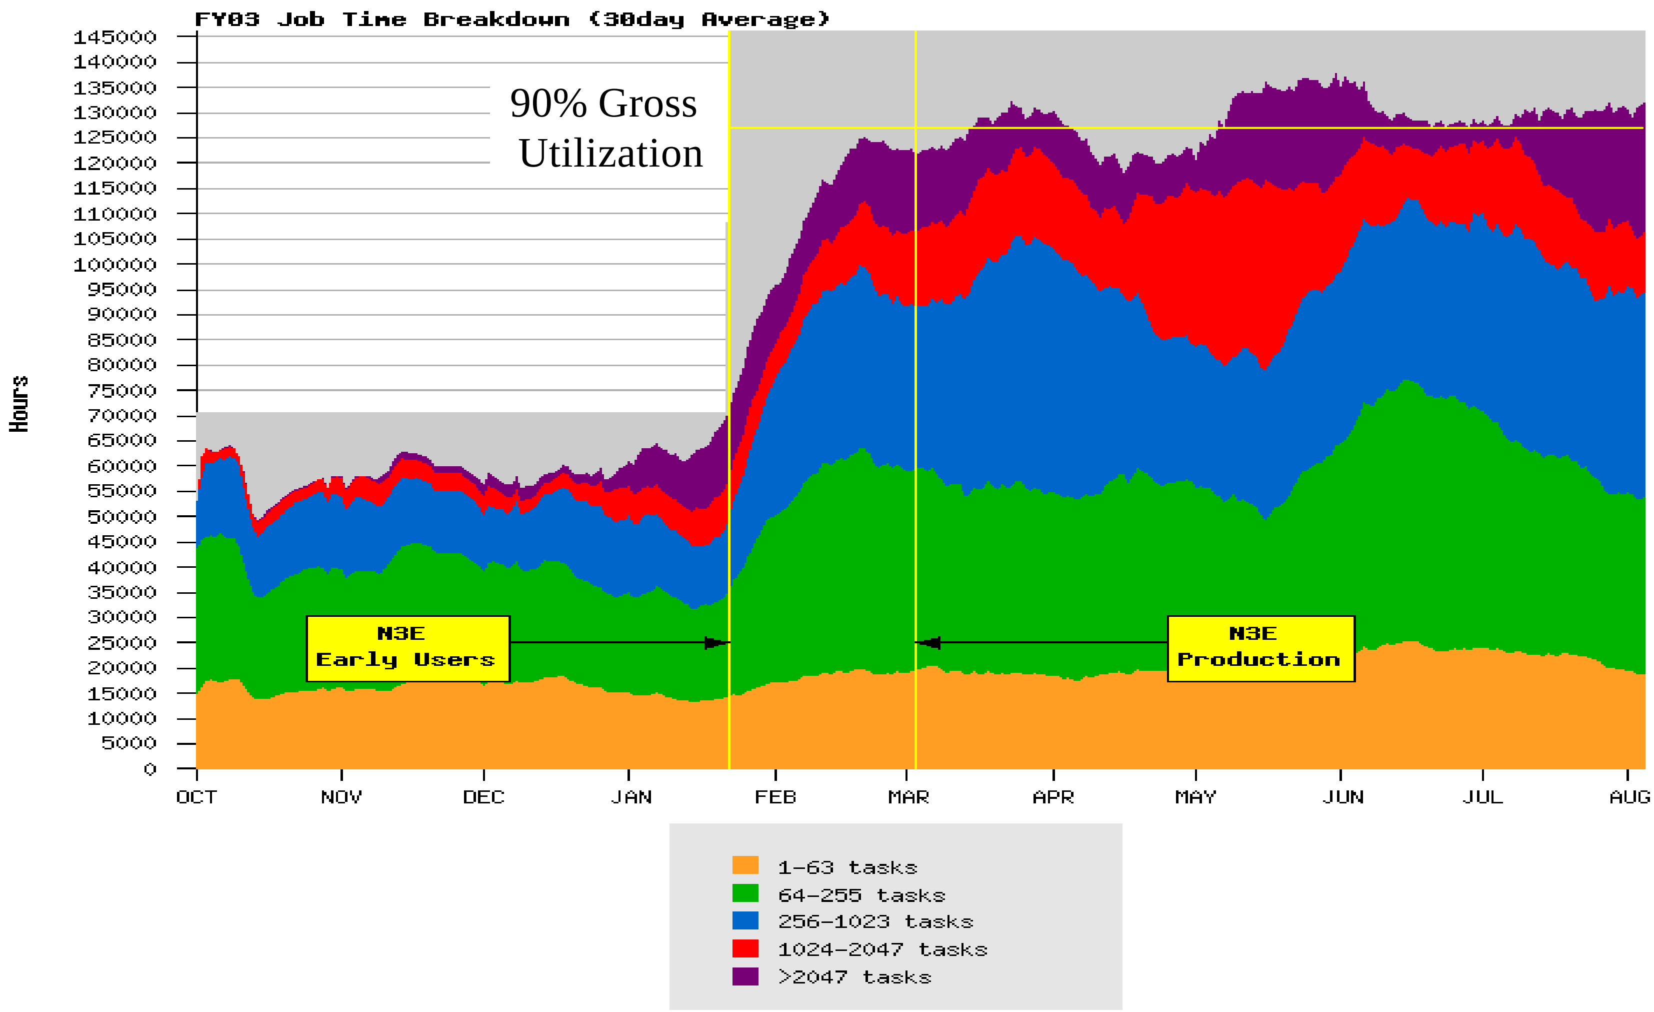

Cumulative Storage Use at NERSC

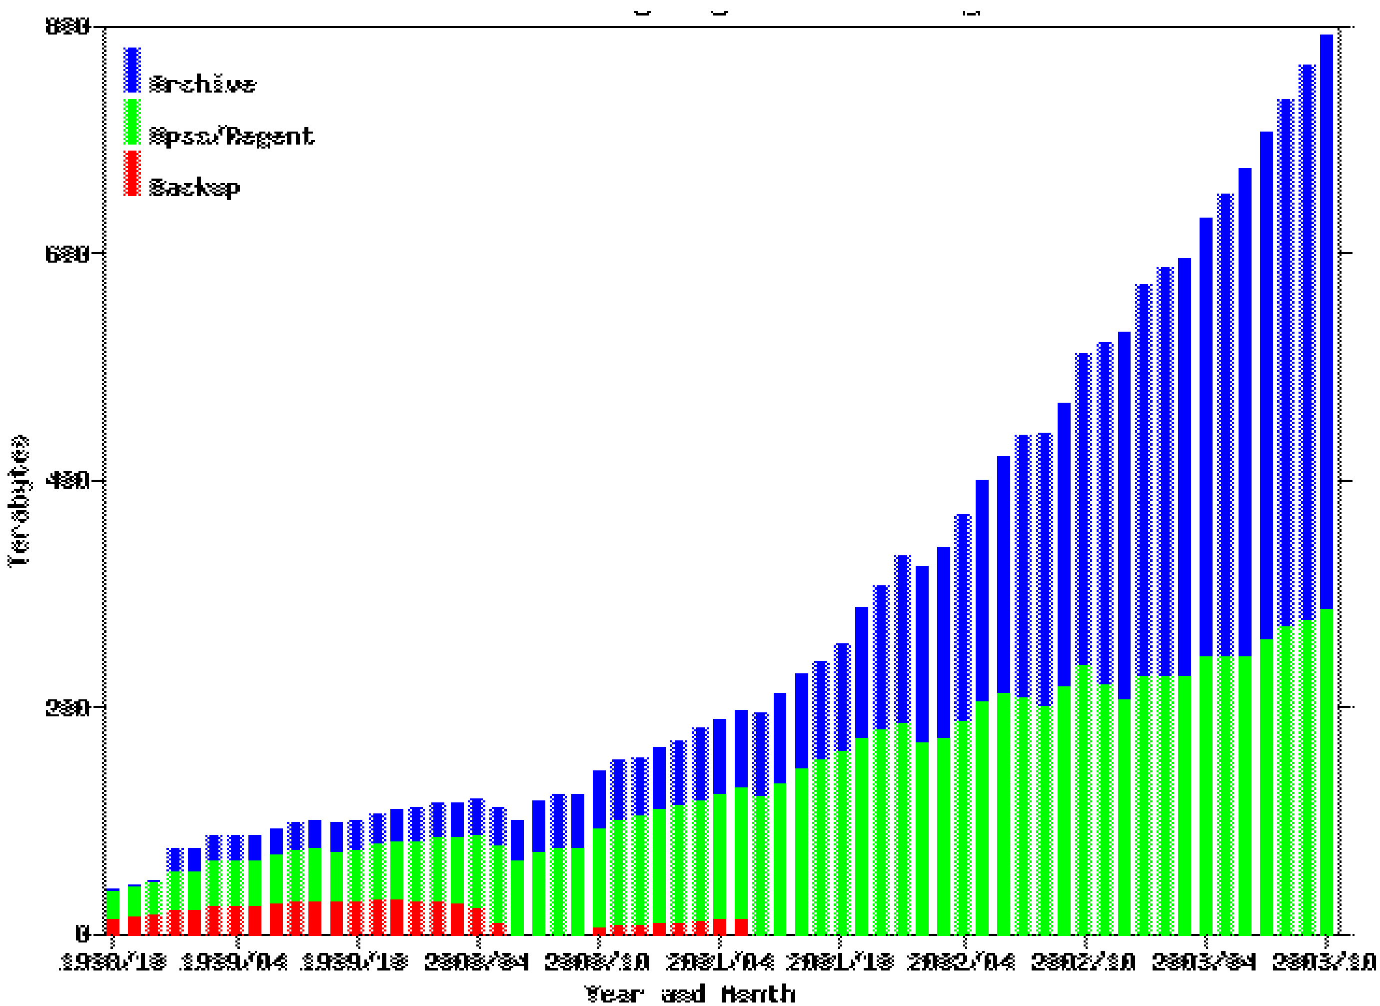

Office of Science
U.S. DEPARTMENT OF ENERGY

Monthly I/O Activity to Storage by Platform = 57 TB

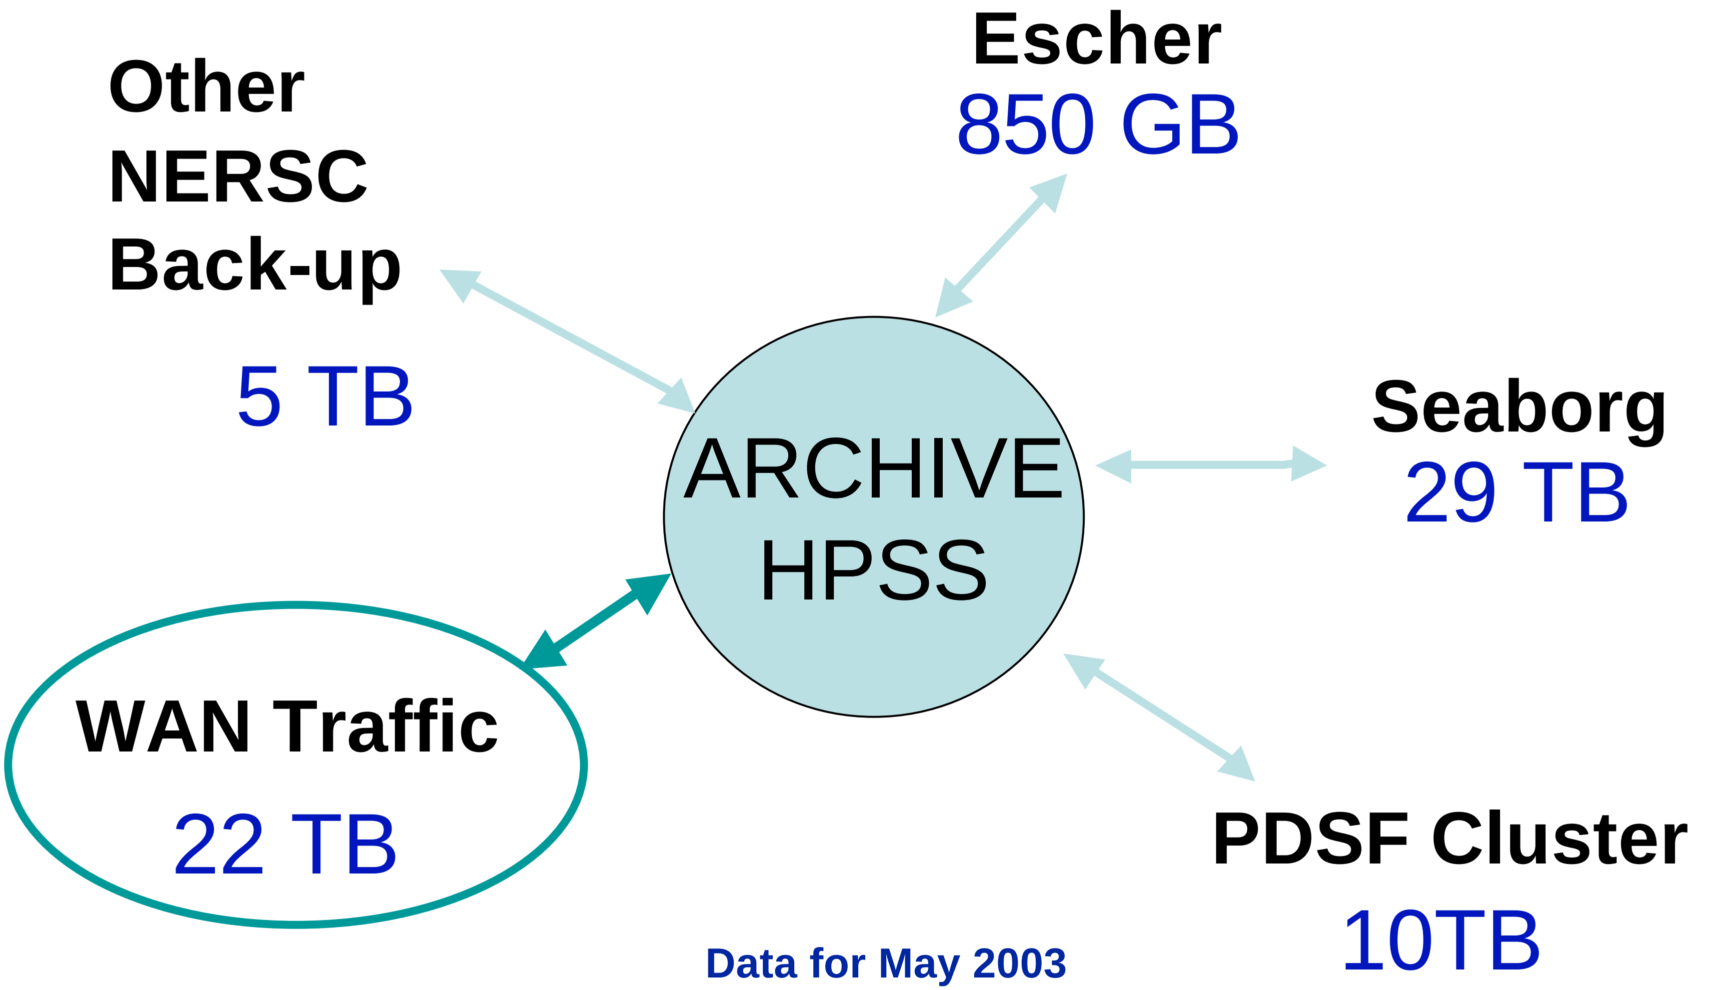

Data for May 2003

NERSC is a Net Importer of Traffic

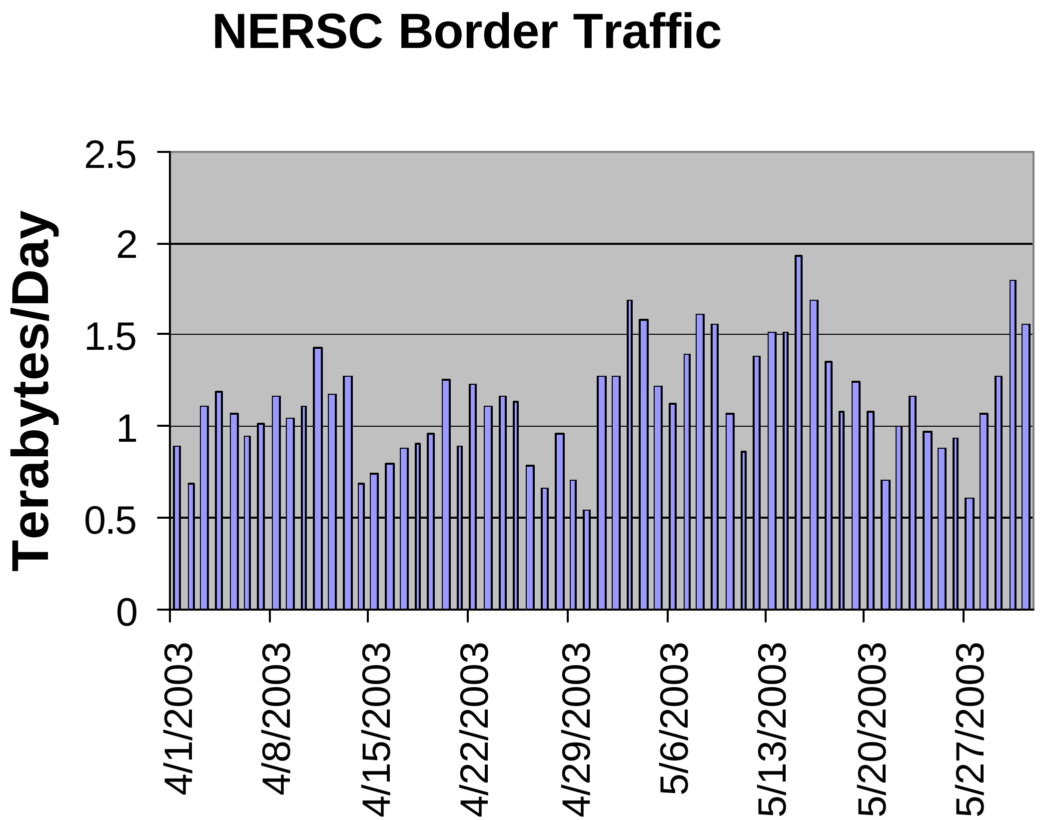

- **Traffic across the NERSC border:**
  - **April 2003 - 29.5 TB**
  - **May 2003 - 39.4 TB**
- **NERSC traffic accounts for approximately 20% of total ESNet traffic**
- **76% of the NERSC traffic is inbound**

Multi-Teraflops Spin Dynamics Studies of the Magnetic Structure of FeMn and FeMn/Co Interfaces

Exchange bias, which involves the use of an antiferromagnetic (AFM) layer such as FeMn to pin the orientation of the magnetic moment of a proximate ferromagnetic (FM) layer such as Co, is of fundamental importance in magnetic multilayer storage and read head devices.

A larger simulation of 4000 atoms of FeMn ran at 4.42 Teraflops on 250 nodes.

(ORNL, Univ. of Tennessee, LBNL(NERSC) and PSC)

IPDPS03 A. Canning, B. Ujfalussy, T.C. Shulthess, X.-G. Zhang, W.A. Shelton, D.M.C. Nicholson, G.M. Stocks, Y. Wang, T. Dirks

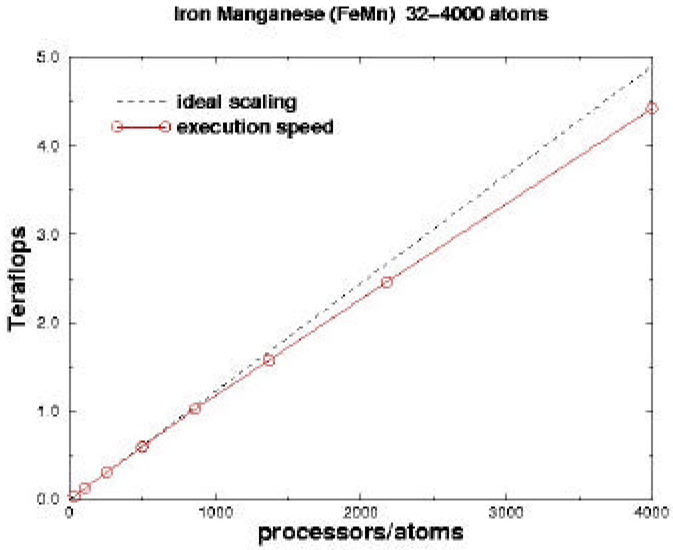

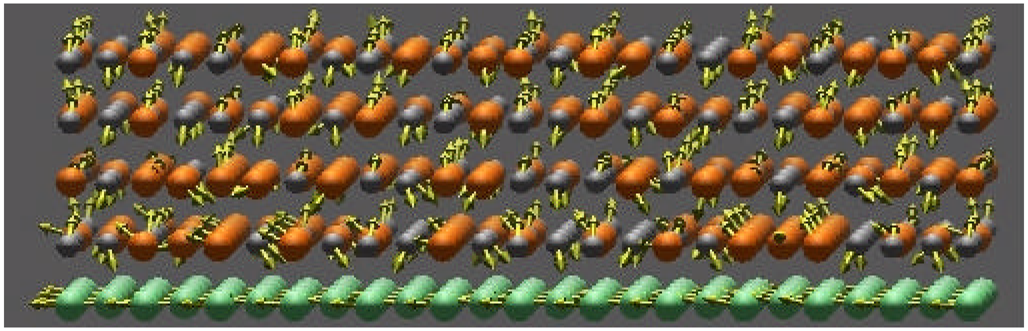

Section of an FeMn/Co (Iron Manganese/ Cobalt) interface showing the final configuration of the magnetic moments for five layers at the interface.

Shows a new magnetic structure which is different from the 3Q magnetic structure of pure FeMn.

NERSC Supports Efficient Science of Scale

Project	Performance (% of peak)	CPU Count
Terascale Simulations of Supernovae	35%	2048
Accelerator Science and Simulation	25%	4096
Electromagnetic Wave-Plasma Interactions	68%	2048
Quantum Chromodynamics at High Temperature	13%	1024

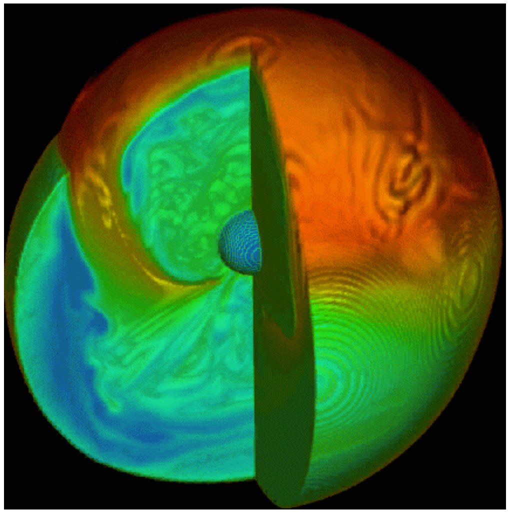

These are real applications, for example:

- **Terascale Simulation of Supernovae**
- PI: **Tony Mezzacappa, ORNL**
- Code**: neutrino scattering on lattices (OAK3D)**
- Kernel**: complex linear equations**
- Performance**: 537 Mflop/s per proc.**
- Scalability**: 1.1 Tflop/s on 2,048 processors**
- Allocation**: 565,000 MPP hours; requested and needs 1.52 million**

NERSC Systems Will Evolve

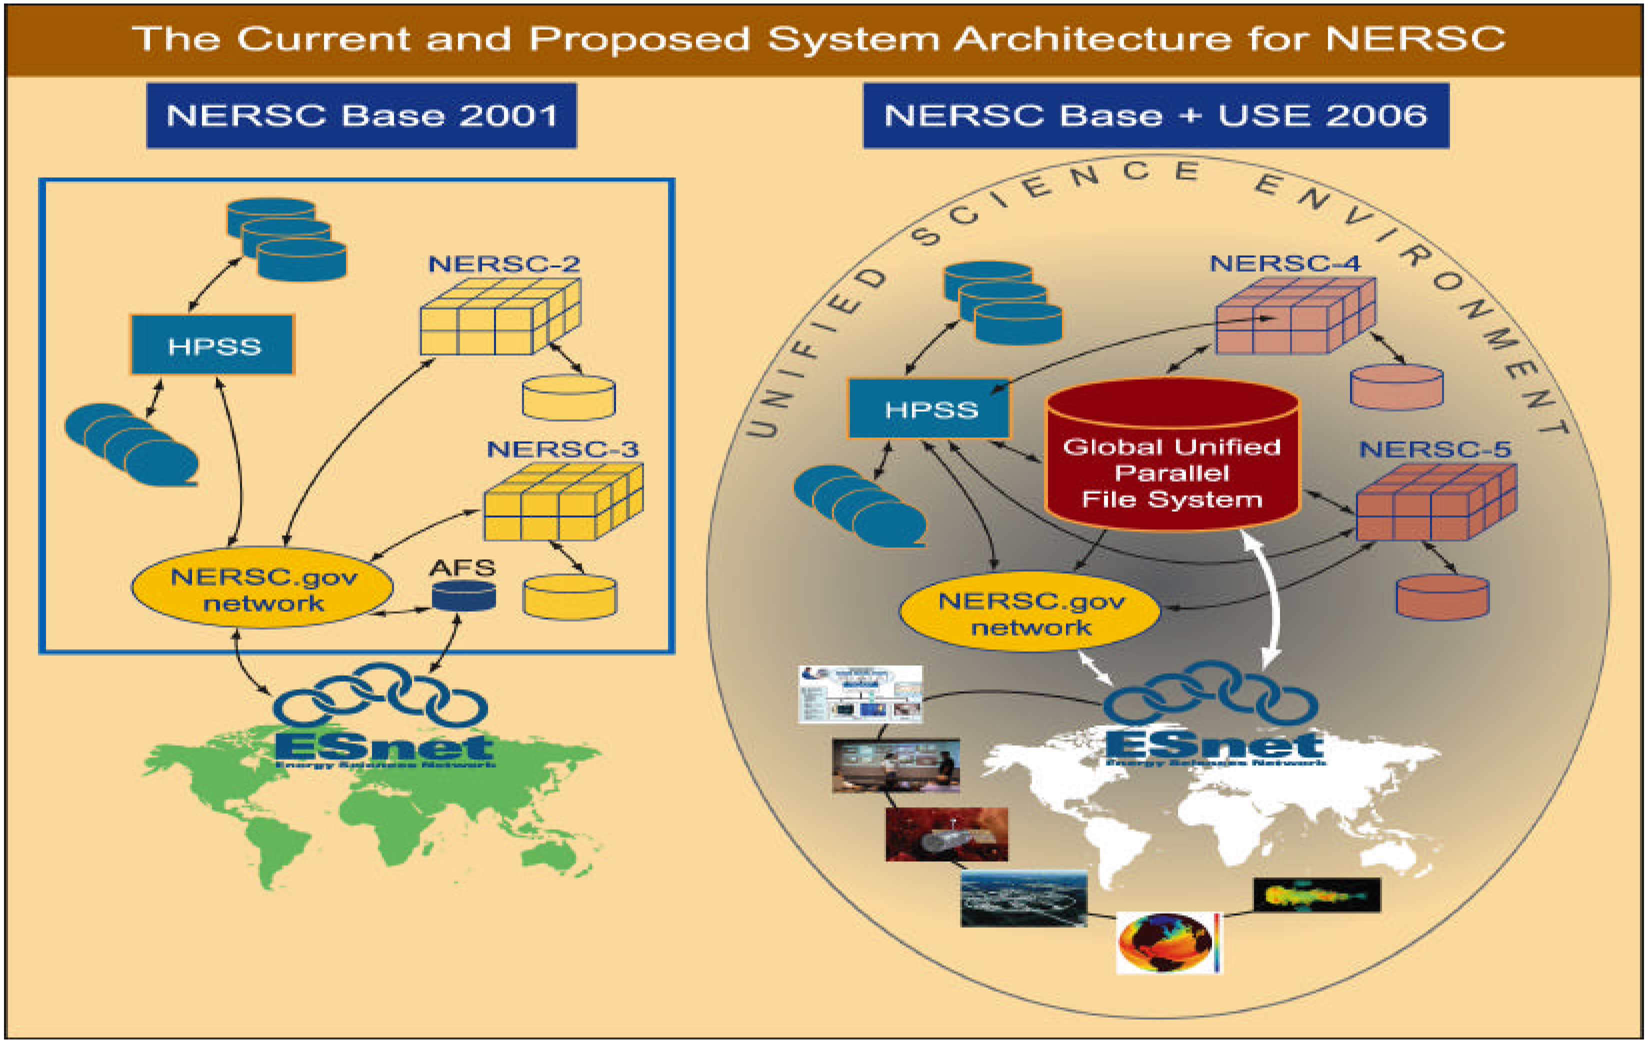

NERSC and the Grid

- Data grid activities are most valuable for NERSC user community
  - PDSF has been grid enabled for about 2 years
  - Developed gridftp for HPSS
  - Collaboration with IBM about deploying grid technology
- Fully subscribed resources at NERSC make computational resource sharing less interesting
- "Nearby Supernova Factory" project at NERSC a prototype example of cyberinfrastructure

Nearby Supernova Factory

- Goal: Find and examine in detail up to 300 nearby Type Ia supernovae
  - More detailed sample against which older, distant supernovae can be compared
- Discovered 34 supernovae during first year of operation and now discovering 8-9 per month
- First year: processed 250,000 images, archived 6 TB of compressed data
- This discovery rate is made possible by:
  - high-speed data link
  - custom data pipeline software
  - NERSC's ability to store and process 50 gigabytes of data every night

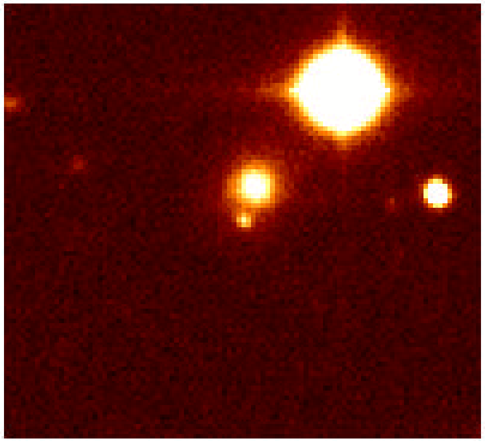

Nearby Supernova Factory

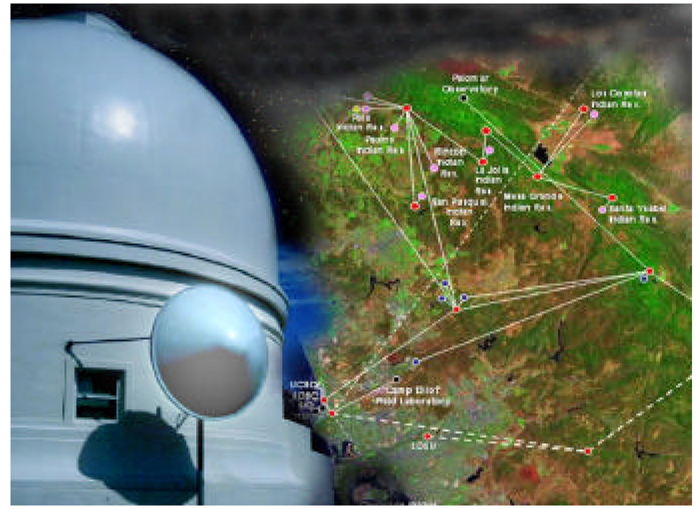

- Every night, images from the Near Earth Asteroid Tracking program (NEAT) at Mount Palomar and Maui are sent to NERSC via ESnet and a special link in SDSC's High Performance Wireless Research and Education Network (HPWREN)
- Custom data pipeline software automatically archives images in NERSC's HPSS
- Image subtraction software running on PDSF sifts through billions of objects to find supernovae
- Follow-up spectrographic observations are obtained the next night and sent to NERSC and other centers for analysis
- First major discovery: First detection of hydrogen in the form of circumstellar material around a supernova

Office of Science
U.S. DEPARTMENT OF ENERGY

Summary

- NERSC provides highly effective computational services and resources in order the meet the needs of DOE Office of Science
  - Highest quality services
  - Extremely usable and effective systems
- NERSC enables computational science not possible elsewhere